A Positive Systems Model of TCP-Like Congestion Control: Asymptotic Results

Robert Shorten, Fabian Wirth, and Douglas Leith

***Abstract*—We study communication networks that employ drop-tail queueing and *Additive-Increase Multiplicative-Decrease (AIMD)* congestion control algorithms. It is shown that the theory of nonnegative matrices may be employed to model such networks. In particular, important network properties, such as: 1) fairness; 2) rate of convergence; and 3) throughput, can be characterized by certain nonnegative matrices. We demonstrate that these results can be used to develop tools for analyzing the behavior of AIMD communication networks. The accuracy of the models is demonstrated by several *NS* studies.**

***Index Terms*—AIMD, congestion control, positive matrices, TCP.**

I. Introduction

In this paper, we describe a design oriented modelling approach that captures the essential features of networks of *Additive-Increase Multiplicative Decrease (AIMD)* sources that employ drop-tail queues. The novelty of our approach lies in the fact that we are able to use the theory of nonnegative matrices and hybrid systems to build mathematical models of communication networks that capture the dynamic interaction between competing flows. This approach is based upon a number of simple observations: 1) communication networks employing congestion control systems are feedback systems; 2) communication systems exhibit event driven phenomena and may therefore be viewed as classical hybrid systems; and 3) network states (queue length, window size, etc.) take only nonnegative values. We show that it is possible to relate important network properties to the characteristics of the nonnegative matrices that arise in the study of such communication networks. In particular, we will demonstrate that: 1) bandwidth allocation amongst flows; 2) rate of network convergence; and 3) network throughput can all be related to properties of sets of nonnegative matrices.

This paper is structured as follows. In Section II, we develop a positive systems network model that captures the essential features of communication networks employing drop-tail queueing and AIMD congestion control algorithms. An exact model is presented for the case where all network sources share a uniform round-trip time (RTT) and packet drops are synchronized. This model is then extended to the case of sources with differing RTTs and where packet drops need not be synchronized. This approach gives rise to a model in which the network dynamics are described by a finite set of nonnegative matrices. The main results of this paper are presented in Section III. To ease exposition these results, which concern the short and long term behavior of AIMD networks, are simply stated in this section. The use of these results to analyze network behavior is illustrated by a number of case studies in Section IV. Finally, in Section VI, we present an outline of the proofs of the mathematical results as well as a number of intermediate derivations. For reasons of space, complete proofs have been transferred to [8].

Manuscript received April 13, 2004; revised January 18, 2005; approved by IEEE/ACM Transactions on Networking Editor F. Paganini. This work was supported by Science Foundation Ireland under grants 00/PI.1/C067 and 04/IN1/I478 and by the Collaborative Research Center 637 "Autonomous Logistic Processes—A Paradigm Shift and its Limitations" funded by the German Research Foundation.

The authors are with the Hamilton Institute, NUI Maynooth, Maynooth, Co. Kildare, Ireland.

Digital Object Identifier 10.1109/TNET.2006.876178

II. Nonnegative Matrices and Communication Networks

A communication network consists of a number of sources and sinks connected together via links and routers. In this paper, we assume that these links can be modelled as a constant propagation delay together with a queue, that the queue is operating according to a drop-tail discipline, and that all of the sources are operating an AIMD-like congestion control algorithm. AIMD congestion control operates a window based congestion control strategy. Each source maintains an internal variable $cwnd_i$ (the window size) which tracks the number of sent unacknowledged packets that can be in transit at any time, i.e., the number of packets in flight. On safe receipt of data packets the destination sends acknowledgement (ACK) packets to inform the source. When the window size is exhausted, the source must wait for an ACK before sending a new packet. Congestion control is achieved by dynamically adapting the window size according to an additive-increase multiplicative-decrease law. Roughly speaking, the idea is for a source to probe the network for spare capacity by increasing the rate at which packets are inserted into the network, and to rapidly decrease the number of packets transmitted through the network when congestion is detected through the loss of data packets. In more detail, the source increments $cwnd_i(t)$ by a fixed amount α_i upon receipt of each ACK. On detecting packet loss, the variable $cwnd_i(t)$ is reduced in multiplicative fashion to $\beta_i cwnd_i(t)$. We shall see that the AIMD paradigm with drop-tail queueing gives rise to networks whose dynamics can be accurately modelled as a positive linear system. While we are ultimately interested in general communication networks, for reasons of exposition it is useful to begin our discussion with a description of networks in which packet drops are synchronized (i.e., every source sees a drop at each congestion event). We show that many of the properties of communication networks that are of interest to network designers can be characterized by properties of a square matrix whose dimension is equal to the number of

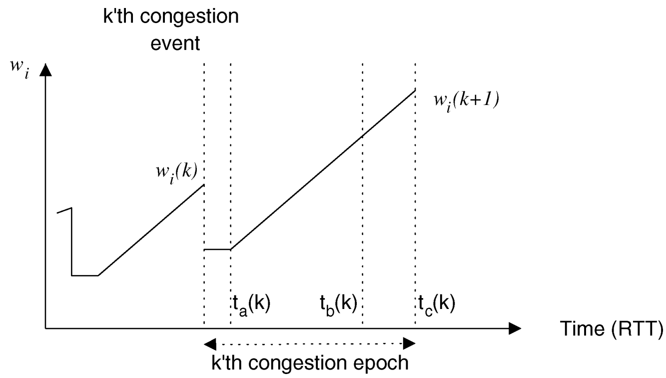

Fig. 1. Evolution of window size.

sources in the network. The approach is then extended to a model of unsynchronized networks. Even though the mathematical details are more involved, many of the qualitative characteristics of synchronized networks carry over to the nonsynchronized case if interpreted in a stochastic fashion.

A. Synchronized Communication Networks

We begin our discussion by considering communication networks for which the following assumptions are valid: 1) at congestion every source experiences a packet drop; and 2) each source has the same RTT.[1] In this case an exact model of the network dynamics may be found as follows [1]. Let $w_i(k)$ denote the congestion window size of source i immediately before the kth network congestion event is detected by the source. Over the kth congestion epoch, three important events can be discerned: $t_a(k)$, $t_b(k)$, and $t_c(k)$, as depicted in Fig. 1. The time $t_a(k)$ denotes the instant at which the number of unacknowledged packets in flight equals $\beta_i w_i(k)$ where β_i is the multiplicative decrease factor associated with the ith flow (recall that after each congestion event the ith sources decreases its number of packets in flight by a factor of $1-\beta_i$); $t_b(k)$ is the time at which the bottleneck queue is full; and $t_c(k)$ is the time at which packet drop is detected by the sources, where time is measured in units of RTT.[2] It follows from the definition of the AIMD algorithm that the window evolution is completely defined over all time instants by knowledge of the $w_i(k)$ and the event times $t_a(k)$, $t_b(k)$, and $t_c(k)$ of each congestion epoch. We therefore only need to investigate the behavior of these quantities.

We assume that each source is informed of congestion one RTT after the queue at the bottleneck link becomes full, that is $t_c(k) - t_b(k) = 1$. Also

$$w_i(k) \geq 0, \sum_{i=1}^{n} w_i(k) = P + \sum_{i=1}^{n} \alpha_i \; \forall k > 0 \tag{1}$$

where P is the maximum number of packets which can be in transit in the network at any time; P is usually equal to $q_{max} + BT_d$ where q_{max} is the maximum queue length of the congested link, B is the service rate of the congested link in packets per second and T_d is the round-trip time when the queue is empty. At the $(k+1)$th congestion event

$$w_i(k+1) = \beta_i w_i(k) + \alpha_i \left[t_c(k) - t_a(k)\right]. \tag{2}$$

It follows from (1) and (2) that

$$t_c(k) - t_a(k) = \frac{1}{\sum_{i=1}^{n} \alpha_i} \left[P - \sum_{i=1}^{n} \beta_i w_i(k)\right] + 1. \tag{3}$$

Hence, it follows that

$$w_i(k+1) = \beta_i w_i(k) + \frac{\alpha_i}{\sum_{j=1}^{n} \alpha_j} \left[\sum_{j=1}^{n} (1-\beta_j) w_j(k)\right] \tag{4}$$

and that the dynamics of an entire network of such sources is given by

$$W(k+1) = AW(k) \tag{5}$$

where $W^T(k) = [w_1(k), \cdots, w_n(k)]$, and where with $\alpha^T = [\alpha_1 \dots \alpha_n]$ we have

$$A = \text{diag}(\beta_1, \beta_2, \dots, \beta_n) + \frac{1}{\sum_{j=1}^{n} \alpha_j} \alpha [1-\beta_1 \; \cdots \; 1-\beta_n] \tag{6}$$

and the initial condition $W(0)$ is subject to the constraint (1).

The matrix A is a positive matrix (all the entries are positive real numbers) and it follows that the synchronized network (5) is a positive linear system [2]. Many results are known for positive matrices and we exploit some of these to characterize the properties of synchronized communication networks. In particular, from the viewpoint of designing communication networks, the following properties are important: 1) network fairness; 2) network convergence and responsiveness; and 3) network throughput. While there are many interpretations of network fairness, in this paper we concentrate on window fairness. Roughly speaking, window or pipe fairness refers to a steady state situation where n sources operating AIMD algorithms have an equal number of packets P/n in flight at each congestion event; convergence refers to the existence of a unique fixed point to which the network dynamics converge; responsiveness refers to the rate at which the network converges to the fixed point; and throughput efficiency refers to the objective that the network operates at close to the bottleneck-link capacity. It is shown in [3] and [4] that these properties can be deduced from the network matrix A. We briefly summarize here the relevant results in these papers.

Theorem 1: [1], [4] Let A be defined as in (6). Then A is a column stochastic matrix with Perron eigenvector $x_p^T = [\alpha_1/(1-\beta_1), \dots, \alpha_n/(1-\beta_n)]$ and whose eigenvalues are real and positive. Further, the network converges to a unique stationary point $W_{ss} = \Theta x_p$, where Θ is a positive constant such that the constraint (1) is satisfied; $\lim_{k\to\infty} W(k) = W_{ss}$; and the rate of convergence of the network to W_{ss} is bounded by the second largest eigenvalue of A.

The following results may be deduced from the above.

[1]One RTT is the time between sending a packet and receiving the corresponding acknowledgement when there are no packet drops.

[2]Note that measuring time in units of RTT results in a linear rate of increase for each of the congestion window variables between congestion events.

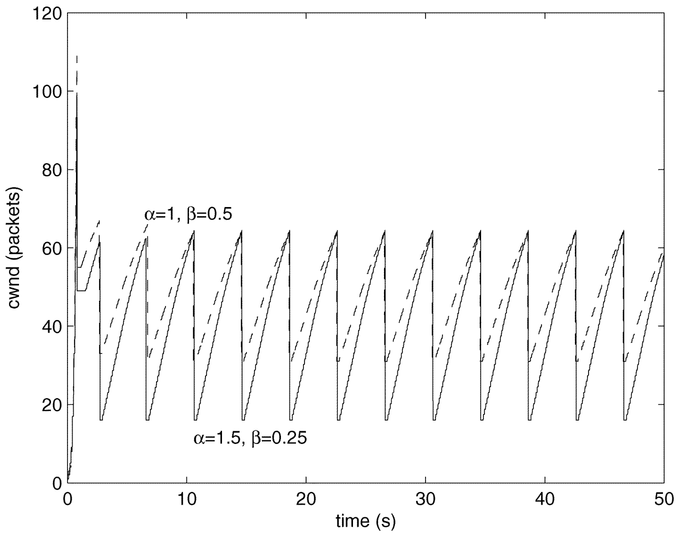

Fig. 2. Example of window fairness between two TCP sources with different increase and decrease parameters (*NS* simulation, network parameters: 10 Mb bottleneck link, 100 ms delay, queue 40 packets.)

(i) Fairness: Window fairness is achieved when the Perron eigenvector x_p is a scalar multiple of the vector $[1, \ldots, 1]$; that is, when the ratio $\alpha_i/(1-\beta_i)$ does not depend on i. Further, since it follows for conventional TCP-flows ($\alpha = 1, \beta = 1/2$) that $\alpha = 2(1-\beta)$, any new protocol operating an AIMD variant that satisfies $\alpha_i = 2(1-\beta_i)$ will be TCP-friendly—i.e., fair with legacy TCP flows.

(ii) Network responsiveness: The magnitude of the second largest eigenvalue λ_{n-1} of the matrix A bounds the convergence properties of the entire network. It is shown in [4] that all the eigenvalues of A are real and positive and lie in the interval $[\beta_1, 1]$, where the β_i are ordered as $0 < \beta_1 \leq \beta_2 \leq \ldots \leq \beta_{n-1} \leq \beta_n < 1$. In particular, the second largest eigenvalue is bounded by $\beta_{n-1} \leq \lambda_{n-1} \leq \beta_n$. Consequently, fast convergence to the equilibrium state (the Perron eigenvector) is guaranteed if the largest backoff factor in the network is small. Further, we show in [4] that the network rise-time when measured in number of congestion epochs is bounded by $n_r = \log(0.95)/\log(\lambda_{n-1})$. In the special case when $\beta_i = 0.5$ for all i, $n_r \approx 4$; see, for example, Fig. 3. Note that n_r gives the number of congestion epochs until the network dynamics have converged to 95% of the final network state: the actual time to reach this state depends on the duration of the congestion epochs which is ultimately dependent on the α_i.

(iii) Network throughput: At a congestion event the network bottleneck is operating at link capacity and the total data throughput through the bottleneck link is given by

$$R(k)^- = \frac{\sum_i^n w_i(k)}{T_d + \frac{q_{max}}{B}} \tag{7}$$

where B is the link capacity, q_{max} is the bottleneck buffer size, T_d is the round-trip time when the bottleneck queue is empty

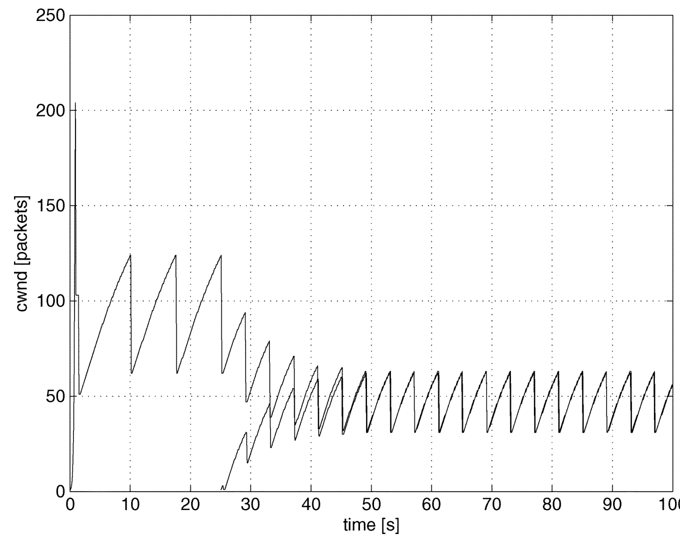

Fig. 3. *NS* packet-level simulation ($\alpha_i = 1, \beta_i = 0.5$, dumbbell with 10 Mbs bottleneck bandwidth, 100 ms propagation delay, 40 packet queue).

and $T_d + q_{max}/B$ is the round-trip time when the queue is full. After backoff, the data throughput is given by

$$R(k)^+ = \frac{\sum_i^n \beta_i w_i(k)}{T_d} \tag{8}$$

under the assumption that the bottleneck buffer empties. It is evident that if the sources backoff too much, data throughput will suffer as the queue remains empty for a period of time and the link operates below its maximum rate. A simple method to ensure maximum throughput is to equate both rates, which may be achieved by the following choice of the β_i:

$$\beta_i = \frac{T_d}{T_d + \frac{q_{max}}{B}} = \frac{\mathrm{RTT}_{min}}{\mathrm{RTT}_{max}}. \tag{9}$$

(iv) Maintaining fairness: Note that setting $\beta_i = \mathrm{RTT}_{min}/\mathrm{RTT}_{max}$ requires a corresponding adjustment of α_i if it is not to result in unfairness. Both network fairness and TCP-friendliness are ensured by adjusting α_i according to $\alpha_i = 2(1-\beta_i)$.

B. Models of Unsynchronized Network

The preceding discussion illustrates the relationship between important network properties and the eigensystem of a positive matrix. Unfortunately, the assumptions under which these results are derived, namely of source synchronization and uniform RTT, are quite restrictive (although they may, for example, be valid in many long-distance networks [5]). It is therefore of great interest to extend our approach to more general network conditions. To distinguish variables, we will from now on denote the nominal parameters of the sources used in the previous section by $\alpha_i^s, \beta_i^s, i = 1, \ldots, n$.

Consider the general case of a number of sources competing for shared bandwidth in a generic dumbbell topology (where sources may have different round-trip times and drops need not

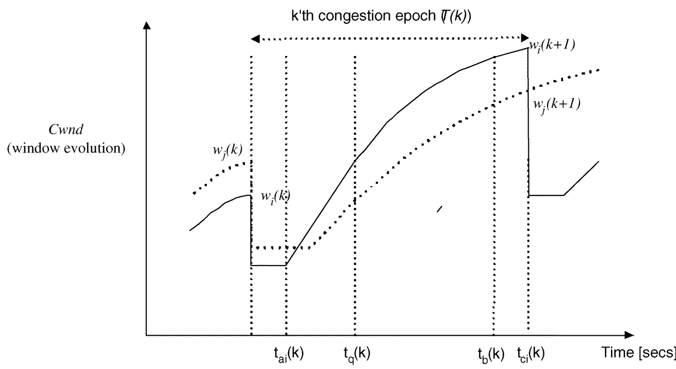

Fig. 4. Evolution of window size over a congestion epoch. $T(k)$ is the length of the congestion epoch in seconds.

be synchronized). The evolution of the $cwnd$ of a typical source as a function of time, over the kth congestion epoch, is depicted in Fig. 4. As before a number of important events may be discerned, where we now measure time in seconds, rather than units of RTT. Denote by $t_{ai}(k)$ the time at which the number of packets in flight belonging to source i is equal to $\beta_i^s w_i(k)$; $t_q(k)$ is the time at which the bottleneck queue begins to fill;[3] $t_b(k)$ is the time at which the bottleneck queue is full; and $t_{ci}(k)$ is the time at which the ith source is informed of congestion. In this case the evolution of the ith congestion window variable does not evolve linearly with time after t_q seconds due to the effect of the bottleneck queue filling and the resulting variation in RTT; namely, the RTT of the ith source increases according to $\mathrm{RTT}_i(t) = T_{d_i} + q(t)/B$ after t_q, where T_{d_i} is the RTT of source i when the bottleneck queue is empty and $0 \le q(t) \le q_{max}$ denotes the number of packets in the queue. Note also that we do not assume that every source experiences a drop when congestion occurs. For example, a situation is depicted in Fig. 4 where the ith source experiences congestion at the end of the epoch, whereas the jth source does not.

Given these general features, it is clear that the modelling task is more involved than in the synchronized case. Nonetheless, it is possible to relate $w_i(k)$ and $w_i(k+1)$ using a similar approach to the synchronized case by accounting for the effect of nonuniform RTTs and unsynchronized packet drops as follows.

(i) Unsynchronized source drops: Consider again the situation depicted in Fig. 4. Here, the ith source experiences congestion at the end of the epoch whereas the jth source does not. This corresponds to the ith source reducing its congestion window by the factor β_i^s after the $(k+1)$th congestion event, and the jth source not adjusting its window size at congestion. We therefore allow the backoff factor of the ith source to take one of two values at the kth congestion event corresponding to whether the source experiences a packet loss or not:

$$\beta_i(k) \in \{\beta_i^s, 1\}. \tag{10}$$

(ii) Nonuniform RTT: Due to the variation in round trip time, the congestion window of a flow does not evolve linearly with time over a congestion epoch. Nevertheless, we may relate $w_i(k)$ and $w_i(k+1)$ linearly by defining an average rate $\alpha_i(k)$ depending on the kth congestion epoch:

$$\alpha_i(k) := \frac{w_i(k+1) - \beta_i(k)w(k)}{T(k)} \tag{11}$$

where $T(k)$ is the duration of the kth epoch measured in seconds. Equivalently, we have

$$w_i(k+1) = \beta_i(k)w_i(k) + \alpha_i(k)T(k). \tag{12}$$

In the case when $q_{max} \ll BT_{d_i}$, $i = 1, \ldots, n$, the average α_i are (almost) independent of k and given by $\alpha_i(k) \approx \alpha_i^s / T_{d_i}$ for all $k \in \mathbb{N}$, $i = 1, \ldots, n$. The situation when

$$\alpha_i \approx \frac{\alpha_i^s}{T_{d_i}}, \quad i = 1, \ldots, n \tag{13}$$

is of considerable practical importance and such networks are the principal concern of this paper. This corresponds to the case of a network whose bottleneck buffer is small compared with the delay-bandwidth product for all sources utilizing the congested link. Such conditions prevail on a variety of networks, for example, networks with large delay-bandwidth products, and networks where large jitter and/or latency cannot be tolerated. Note, however, that the model is not restricted to such situations; see Comment 4.1 below. In view of (10) and (12), a convenient representation of the network dynamics is obtained as follows. At congestion, the bottleneck link is operating at its capacity B, i.e.,

$$\sum_{i=1}^{n} \frac{w_i(k) - \alpha_i}{\mathrm{RTT}_{i,max}} = B \tag{14}$$

where $\mathrm{RTT}_{i,max}$ is the RTT experienced by the ith flow when the bottleneck queue is full. Note, that $\mathrm{RTT}_{i,max}$ is independent of k. Setting $\gamma_i := (\mathrm{RTT}_{i,max})^{-1}$ we have that

$$\sum_{i=1}^{n} \gamma_i w_i(k) = B + \sum_{i=1}^{n} \gamma_i \alpha_i. \tag{15}$$

By interpreting (15) at $k+1$ and inserting (12) for $w_i(k+1)$

$$\sum_{i=1}^{n} \gamma_i \beta_i(k) w_i(k) + \gamma_i \alpha_i T(k) = B + \sum_{i=1}^{n} \gamma_i \alpha_i. \tag{16}$$

Using (15) again it follows that

$$T(k) = \frac{1}{\sum_{i=1}^{n} \gamma_i \alpha_i} \left(\sum_{i=1}^{n} \gamma_i \left(1 - \beta_i(k)\right) w_i(k) \right). \tag{17}$$

Inserting this expression into (12) we finally obtain

$$w_i(k+1) = \beta_i(k) w_i(k) + \frac{\alpha_i}{\sum_{j=1}^{n} \gamma_j \alpha_j} \left(\sum_{j=1}^{n} \gamma_j \left(1 - \beta_j(k)\right) w_j(k) \right). \tag{18}$$

and the dynamics of the entire network of sources at the kth congestion event, subject to (15), are described by

$$W(k+1) = A(k)W(k), \quad A(k) \in \{A_1, \ldots, A_m\} \tag{19}$$

[3]In the case when the queue does not empty following a congestion event $t_q(k) = t_{ai}(k)$.

where setting $g^T = [\gamma_1(1-\beta_1(k)) \, \ldots \, \gamma_n(1-\beta_n(k))]$ we have

$$A(k) = \mathrm{diag}\left(\beta_1(k), \beta_2(k), \ldots, \beta_n(k)\right) + \frac{1}{\sum_{j=1}^{n} \gamma_j \alpha_j} \alpha g^T \tag{20}$$

and where $\beta_i(k)$ is either 1 or β_i^s. The nonnegative matrices $A_2, \ldots, A_m$ are constructed by taking the matrix A_1:

$$A_1 = \begin{bmatrix} \beta_1^s & 0 & \cdots & 0 \\ 0 & \beta_2^s & 0 & 0 \\ \vdots & 0 & \ddots & 0 \\ 0 & 0 & \cdots & \beta_n^s \end{bmatrix} + \frac{1}{\sum_{j=1}^{n} \gamma_j \alpha_j} \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \cdots \\ \alpha_n \end{bmatrix} \times \left[\gamma_1 \left(1-\beta_1^s\right), \ldots, \gamma_n \left(1-\beta_n^s\right)\right] \tag{21}$$

and setting some, but not all, of the β_i^s to 1. This gives rise to a set $\mathcal{A}$ of $m = 2^n - 1$ matrices associated with the system (19) that correspond to the different combinations of source drops that are possible.

Finally, we note that another, sometimes very useful, representation of the network dynamics can be obtained by considering the evolution of scaled window sizes at congestion, namely, by considering the evolution of $W_\gamma^T(k) = [\gamma_1 w_1(k), \gamma_2 w_2(k), \ldots, \gamma_n w_n(k)]$. Here one obtains the following description of the network dynamics:

$$W_\gamma(k+1) = \bar{A}(k) W_\gamma(k), \; \bar{A}(k) \in \bar{\mathcal{A}} = \{\bar{A}_1, \ldots, \bar{A}_m\} \tag{22}$$

where $m = 2^n - 1$ and the $\bar{A}_i$ are obtained by the similarity transformation associated with the change of variables, in particular

$$\bar{A}_1 = \begin{bmatrix} \beta_1^s & 0 & \cdots & 0 \\ 0 & \beta_2^s & 0 & 0 \\ \vdots & 0 & \ddots & 0 \\ 0 & 0 & \cdots & \beta_n^s \end{bmatrix} + \frac{1}{\sum_{j=1}^{n} \gamma_j \alpha_j} \begin{bmatrix} \gamma_1 \alpha_1 \\ \gamma_2 \alpha_2 \\ \cdots \\ \gamma_n \alpha_n \end{bmatrix} \times \left[1-\beta_1^s \quad 1-\beta_2^s \quad \cdots \quad 1-\beta_n^s\right].$$

As before, the nonnegative matrices $\bar{A}_2, \ldots, \bar{A}_m$ are constructed by taking the matrix $\bar{A}_1$ and setting some, but not all, of the β_i^s to 1. All of the matrices in the set $\bar{\mathcal{A}}$ are now column stochastic; in our proofs we use this representation of the network dynamics.

III. Main Results

The ultimate objective of our work is to use the network model developed in Section II to establish design principles for the realisation of AIMD networks. In this section, we present two results, both of which are derived from our network model in Section II, which in a sense characterize the asymptotic behavior of both long and short lived flows.

A. Preamble to Main Results

It follows from (19) that $W(k) = \Pi(k)W(0)$, where $\Pi(k) = A(k)A(k-1)\ldots A(0)$. Consequently, the behavior of $W(k)$, as well as the network fairness and convergence properties, are governed by the properties of the matrix product $\Pi(k)$. The objective of this section is to analyze the average behavior of $\Pi(k)$ with a view to making concrete statements about these network properties. To facilitate analytical tractability, we will make two mild simplifying assumptions.

Assumption 1: The probability that $A(k) = A_i$ in (19) is independent of k and equals ρ_i.

Comment 1: In other words, Assumption 1 says that the probability that the network dynamics are described by $W(k+1) = A(k)W(k)$, $A(k) = A_i$ over the kth congestion epoch is ρ_i and that the random variables $A(k)$, $k \in \mathbb{N}$ are independent and identically distributed (i.i.d.).

Given the probabilities ρ_i for $i \in \{1, \ldots, 2^n - 1\}$, one may then define the probability λ_j that source j experiences a loss event at the kth congestion event as follows:

$$\lambda_j = \sum \rho_i$$

where the summation is taken over those i which correspond to a matrix in which the jth source sees a drop. Or to put it another way, the summation is over those indexes i for which the matrix A_i is defined with a value of $\beta_j \neq 1$. Note that λ_j can be thought of as a synchronization factor—it is unity when a flow experiences a loss at every congestion event.

Assumption 2: Let $\lambda_j > 0$ for all $j \in \{1, \ldots, n\}$.

Simply stated, Assumption 2 states that almost surely all flows must see a drop at some time (provided that they live for a long enough time).

Comment 2: A consequence of the above assumptions is that the probability that source j experiences a drop at the kth congestion event is not independent of the other sources. For example, if the first $n-1$ sources do not see a drop then this implies that source n must see a drop (in accordance with the usual notion of a congestion event, we require at least one flow to see a drop at each congestion event).

We now present two results that characterize the expected behavior of AIMD networks that satisfy Assumptions 1 and 2. The first characterizes the ensemble average behavior of flows, while the second characterizes the time average behavior.

B. Result 1. Ensemble Average Behavior of TCP-Flows

Theorem 2: Consider the stochastic system defined in the preamble. Let $\Pi(k)$ be the random matrix product arising from the evolution of the first k steps of this system:

$$\Pi(k) = A(k)A(k-1)\ldots A(0).$$

Then, the expectation of $\Pi(k)$ is given by

$$E\left(\Pi(k)\right) = \left(\sum_{i=1}^{m} \rho_i A_i\right)^k \tag{23}$$

and the asymptotic behavior of $E(\Pi(k))$ satisfies

$$\lim_{k \to \infty} E\left(\Pi(k)\right) = x_p y_p^T \tag{24}$$

where $x_p^T = \Theta(\alpha_1/\lambda_1(1-\beta_1), \alpha_2/\lambda_2(1-\beta_2), \ldots, \alpha_n/\lambda_n(1-\beta_n))$, $y_p^T = (\gamma_1, \ldots, \gamma_n)$. Here $\Theta \in \mathbb{R}$ is chosen such that (15) is satisfied if w_i is replaced by $x_{pi} = \Theta\alpha_i/(\lambda_i(1-\beta_i))$.

Comment 3: Theorem 2 characterizes the ensemble average behavior of the congestion variable vector $W(k)$. The congestion variable vector of a network of flows starting from initial condition $W(0)$ and evolving for k congestion epochs is given by $W(k) = \Pi(k)W(0)$. The average window vector over many repetitions is given by $E(\Pi(k))W(0)$. Theorem 2 provides an expression for calculating this average in terms of the network parameters and the probabilities ρ_i. Furthermore, we have that as k becomes large $E(\Pi(k))W(0)$ tends asymptotically to $x_p y_p^T W(0)$. The rate of convergence of $E(\Pi(k))W(0)$ to this limiting value is bounded by the second largest eigenvalue of the matrix $\sum_{n=1}^{m} \rho_i A_i$.

Comment 4: Theorem 2 is concerned with the expected behavior of the source congestion windows at the kth congestion epoch. For k sufficiently large the expected throughput before backoff can be approximated as $\sum_{i=1}^{n}(\alpha_i/\lambda_i(1-\beta_i)\mathrm{RTT}_{i,max})$. The expected worst case throughput after backoff (which occurs when the queue is on average empty after backoff) is approximately $\sum_{i=1}^{n}(\alpha_i/\lambda_i(1-\beta_i)\mathrm{RTT}_{i,min})$. An immediate consequence of this observation is that the bottleneck link is guaranteed to be operating at capacity (on average) for k large enough if $\beta_i = \mathrm{RTT}_{i,min}/\mathrm{RTT}_{i,max}$.

C. *Result 2. Time Average Behavior of Flows*

We now present the following theorem which is concerned with networks characterized by long-lived flows.

Theorem 3: Consider the stochastic system defined in the preamble and let

$$\overline{W}(k) := \frac{1}{k+1}\sum_{i=0}^{k} W1) = \left(\frac{1}{k+1}\sum_{i=0}^{k}\Pi1)\right) W(0)$$

where $\Pi(k) = A(k)A(k-1)\dots A(0)$. Then, with probability one

$$\lim_{k\to\infty}\overline{W}(k) = x_p y_p^T W(0) = \left(\sum_{j=1}^{n}\gamma_j w_j(0)\right) x_p \tag{25}$$

where the vectors x_p and y_p are as defined in Theorem 2.

Comment 5: (i) Theorem 3 states that the time-average vector of window sizes almost surely converges asymptotically to a scalar multiple of x_p. Hence, x_p determines the time-averaged relative number of unacknowledged packets in the network from each source at each congestion event.

(ii) In view of Comment 4, it again follows that asymptotically, the time-averaged throughput through the bottleneck link will approach the capacity B for k sufficiently large if $\beta_i = \mathrm{RTT}_{i,min}/\mathrm{RTT}_{i,max}$.

(iii) In the spirit of Theorem 2 one may also consider the expectation of $\overline{W}(k)$: $E(\overline{W}(k))$. Denoting $E(A) = \sum_{i=1}^{m}\rho_i A_i$ it follows from our assumptions that

$$E\left(\overline{W}(k)\right) = \frac{1}{k+1}\left(E(A)^k + \dots + E(A)\right)W(0) \tag{26}$$

$$= M(k)W(0) \tag{27}$$

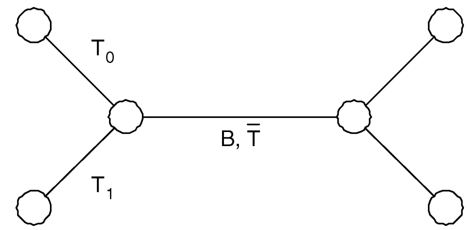

Fig. 5. Dumbbell topology.

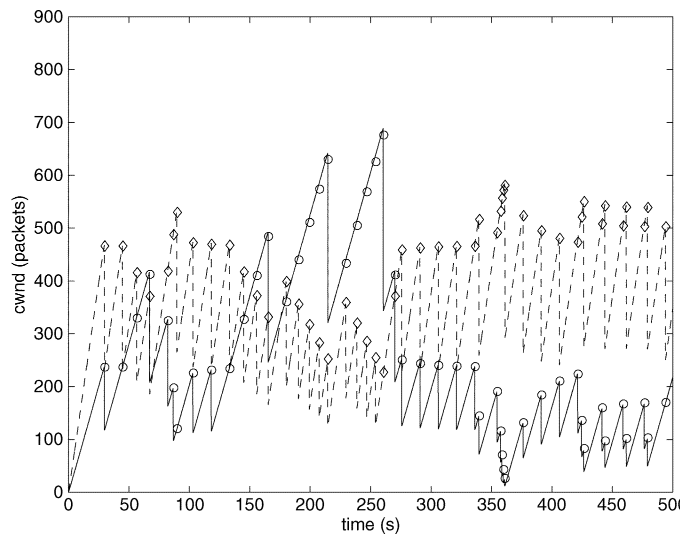

Fig. 6. Predictions of the network model compared with packet-level *NS* simulation results. Key: ∘ flow 1 (model); ◊ flow 2 (model); - flow 1 *(NS)*; — flow 2 *(NS)*. Network parameters: $B = 100$ Mb, $q_{max} = 80$ packets, $\bar{T} = 20$ ms, $T_0 = 102$ ms; $T_1 = 42$ ms; no background web traffic.

where the matrix $M(k)$ is column stochastic matrix, and whose second largest eigenvalue is given by

$$\overline{\lambda_2}(k) = \frac{1}{k+1}\frac{\lambda_2^{k+1}-\lambda_2}{\lambda_2 - 1} \tag{28}$$

where λ_2 is the second largest eigenvalue of $E(A)$. Further, $E(\overline{W}(k))$ tends asymptotically to x_p as $\overline{\lambda_2}(k)$ tends to 0.

IV. Model Validation

The mathematical results derived in Section III are surprisingly simple when one considers the potential mathematical complexity of the unsynchronized network model (19). The simplicity of these results is a direct consequence of Assumptions 1 and 2. The objective of this section is therefore twofold: 1) to validate the unsynchronized model (19) in a general context; and 2) to validate the analytical predictions of the model and thereby confirm that the aforementioned assumptions are appropriate in practical situations.

A. *Two Unsynchronized Flows*

We first consider the behavior of two TCP flows in the dumbbell topology shown in Fig. 5. Our analytic results are based upon two fundamental assumptions: 1) that the dynamics of the evolution of the source congestion windows can be accurately modelled by (19); and 2) the allocation of packet drops amongst

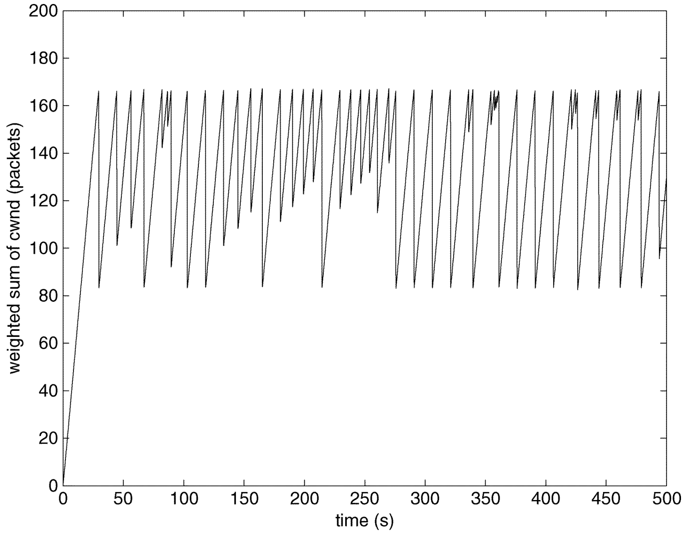

Fig. 7. Evolution of $\sum_{i=1}^{n} \gamma_i w_i$. The peaks correspond to congestion events. Network parameters: $B = 100$ Mb, $q_{max} = 80$ packets, $\bar{T} = 20$ ms, $T_0 = 102$ ms; $T_1 = 42$ ms; no background web traffic.

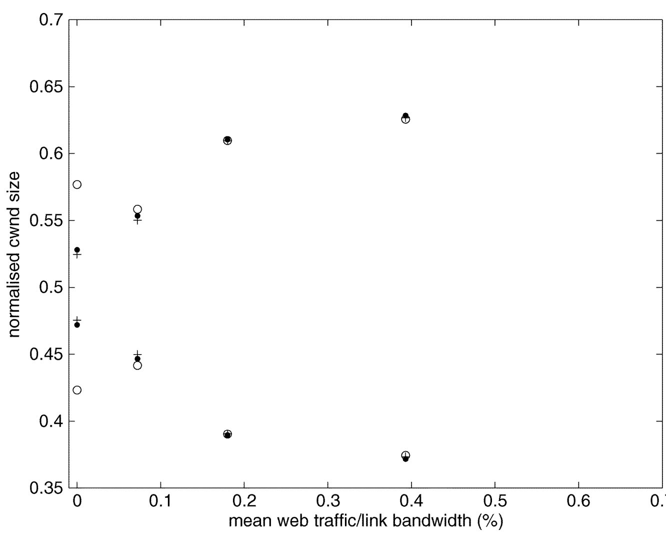

Fig. 8. Variation of mean $w_i(k)$ with level of background web traffic in dumbbell topology of Fig. 5. Key: + *NS* simulation result; · mathematical model (19); ∘ Theorem 3. Network parameters: $B = 100$ Mb, $q_{max} = 80$ packets, $\bar{T} = 20$ ms, $T_0 = 102$ ms; $T_1 = 42$ ms.

the sources at congestion can be described by random variables. We consider each of these assumptions in turn.

1) Accuracy of dynamics model: A comparison of the predictions of the model (19) against the output of a packet-level *NS* simulation is depicted in Fig. 6. Here, the pattern of packet drops observed in the simulation is used to select the appropriate matrix $A(k)$ from the set $\mathcal{A}$ at each congestion event when evaluating (19). As can be seen, the model output is very accurate. Also plotted in Fig. 7 is the evolution of the linear combination $\sum_{i=1}^{n} \gamma_i w_i$ where the γ_i are defined in (15). It can be seen that $\sum_{i=1}^{n} \gamma_i w_i$ has the same value at each congestion event, thereby validating the constraint (15) used in the model.

2) Validity of random drop model: It is well known that networks of TCP flows with drop-tail queues can exhibit a rich variety of deterministic drop behaviors [6]. However, most real

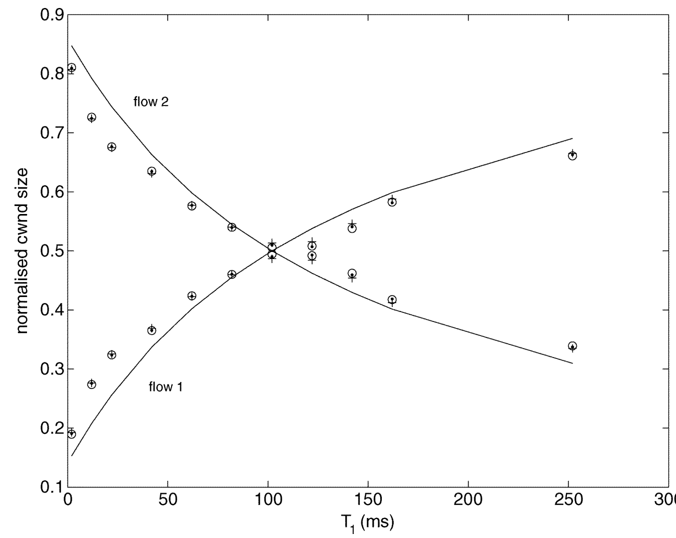

Fig. 9. Variation of mean $w_i(k)$ with propagation delay T_1 in dumbbell topology of Fig. 5. Key: + *NS* simulation result; · mathematical model (19); ∘ Theorem 3; solid lines correspond to synchronized case. Network parameters: $B = 100$ Mb, $q_{max} = 80$ packets, $\bar{T} = 20$ ms, $T_0 = 102$ ms; approximately 0.5% bidirectional background web traffic.

TABLE I
TABULAR DATA FOR FIG. 9

T_1 (ms)	Flow 1		
	NS Simulation	Model	Theorem 3
2.0	0.1924	0.1908	0.1895
12.0	0.2762	0.2757	0.2736
22.0	0.3253	0.3235	0.3237
42.0	0.3691	0.3654	0.3651
62.0	0.4226	0.4230	0.4239
82.0	0.4599	0.4605	0.4600
102.0	0.4866	0.4901	0.4943
122.0	0.5156	0.5071	0.5082
142.0	0.5461	0.5406	0.5378
162.0	0.5877	0.5813	0.5825
252.0	0.6652	0.6627	0.6609
	Flow 2		
	NS Simulation	Model	Theorem 3
2.0	0.8076	0.8092	0.8105
12.0	0.7238	0.7243	0.7264
22.0	0.6747	0.6765	0.6763
42.0	0.6309	0.6346	0.6349
62.0	0.5774	0.5770	0.5761
82.0	0.5401	0.5395	0.5400
102.0	0.5134	0.5099	0.5057
122.0	0.4844	0.4929	0.4918
142.0	0.4539	0.4594	0.4622
162.0	0.4123	0.4187	0.4175
252.0	0.3348	0.3373	0.3391

networks carry at least a small amount of web traffic. In Fig. 8, we plot *NS* simulation results showing the mean congestion window as the level of background web traffic is varied (background information on the web traffic generator in *NS* is described in [7]). To illustrate the impact of small amounts of web traffic, these results are given for a network condition where phase effects are particularly pronounced. While the agreement between the simulation and our random matrix model is poor

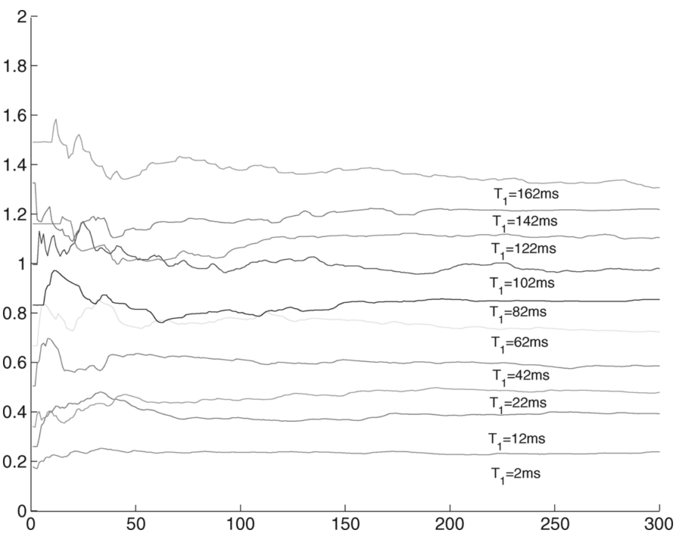

Fig. 10. Convergence of the empirical mean of the window size to asymptotic values shown in Fig. 9. *NS* simulation results; network parameters: $B = 100$ Mb, $q_{max} = 80$ packets, $\bar{T} = 20$ ms, $T_0 = 102$ ms.

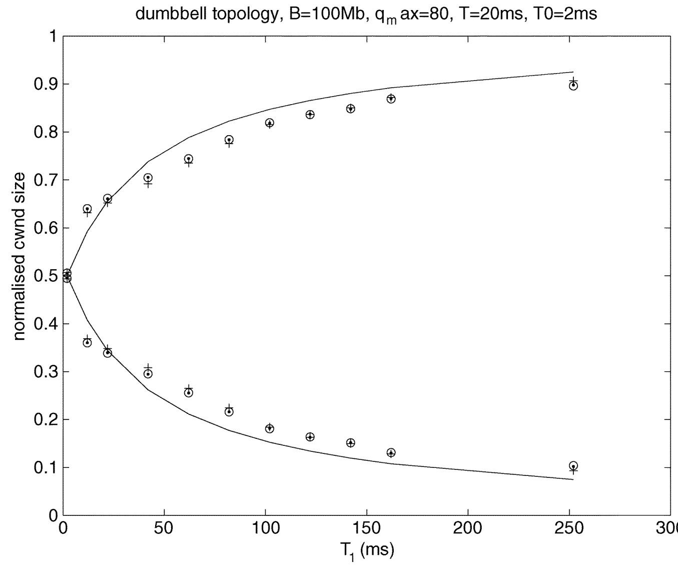

Fig. 11. Variation of mean $w_i(k)$ with propagation delay T_1 in dumbbell topology of Fig. 5. Key: + *NS* simulation result; · mathematical model (19); o Theorem 3; solid lines correspond to synchronized case. Network parameters: $B = 100$ Mb, $q_{max} = 80$ packets, $\bar{T} = 20$ ms, $T_0 = 2$ ms; approximately 0.5% bidirectional background web traffic.

with no web traffic, even a very small volume of web traffic appears to be enough to disrupt the coherent structure associated with phase effects and other complex phenomena previously observed in simulations of unsynchronized networks. From the packet-based simulation results we can determine the proportion of congestion events corresponding to both flows simultaneously seeing a packet drop, flow 1 seeing a drop only, and flow 2 seeing a drop only. Using these estimates of the probabilities ρ_i, the mean congestion window can be estimated using expression (24) from Theorem 3. The resulting estimates are shown in Fig. 9, and are also presented in tabular form in Table I. In these figures T_1 is the fixed delay associated with source 1 that is depicted in Fig. 5. The first column for each flow gives the actual average window size as predicted by the *NS* simulator; the second column gives the predictions of the model (19); and the third column gives the long-time average predictions of Theorem 3. It can be seen that there is close agreement between the packet-level simulation results and the predictions obtained using (24). The actual convergence of the simulation data to the mean values is depicted in Fig. 10.

Also shown in Fig. 9 are the analytic predictions for the case where each source has an equal probability of backing off when congestion occurs, namely, when $\lambda_i = (1/n)\ \forall i$. The corresponding ratio of the elements of the average congestion window vector is the same as that under the assumption of source synchronization (it is important to note that patterns of packet drop other than synchronized drops can lead to the same distribution as long as the proportion of backoff events experienced by the two flows is the same). Observe that the resulting predictions are an accurate estimate of the mean congestion window size and that as the level of web traffic increases the mean window size approaches that in the synchronized case (see Fig. 8).

Before proceeding, we also present results from several other two-flow networks in Figs. 11 and 12. As can be seen from the

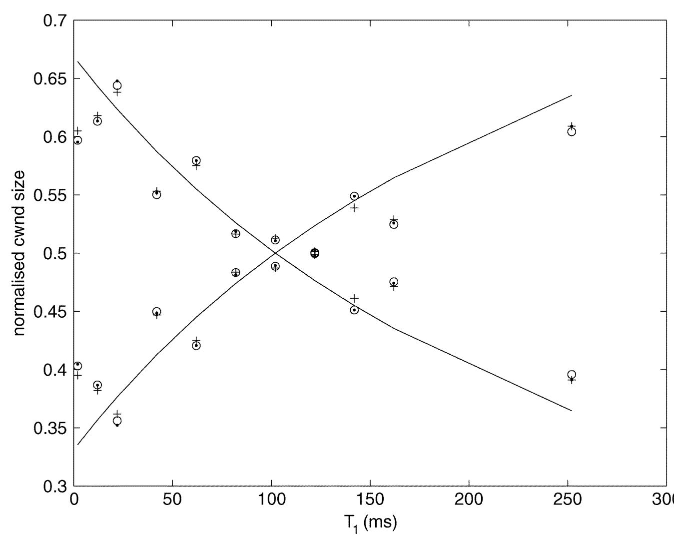

Fig. 12. Variation of mean $w_i(k)$ with propagation delay T_1 in dumbbell topology of Fig. 5. Key: + *NS* simulation result; · mathematical model (19); o Theorem 3; solid lines correspond to synchronized case. Network parameters: $B = 100$ Mb, $q_{max} = 80$ packets, $\bar{T} = 100$ ms, $T_0 = 102$ ms; approximately 0.5% bidirectional background web traffic.

figures, the predictions of Theorem 3 and the *NS* simulations are consistently in close agreement.

The foregoing results are for networks with two competing TCP sources. We note briefly that we have also validated our results against packet-level simulations for networks of up to five flows. As in the two-flow case, the simulation and analytical predictions are in close agreement (to the same degree of accuracy).

Comment 6: (i) Predictions based upon the model (19) rely on knowledge of the rate α_i at which each of the sources increases its window size. In the case of networks with small queue sizes,

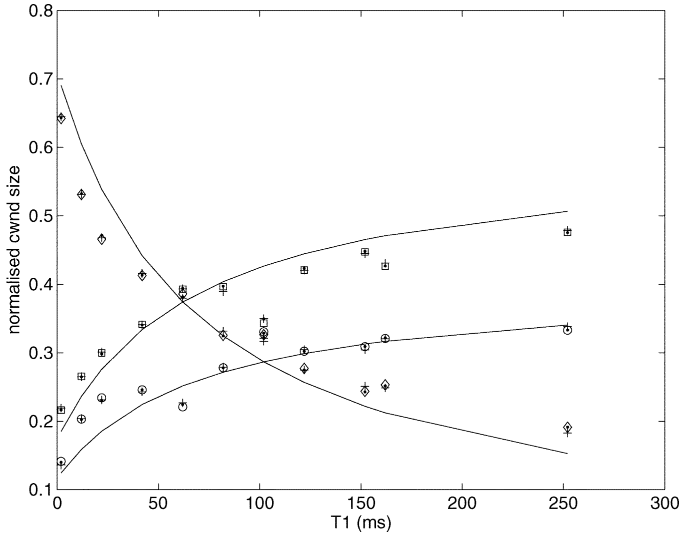

Fig. 13. Variation of mean $w_i(k)$ with propagation delay T_1 in dumbbell topology with three TCP flows. Key: + *NS* simulation result; · mathematical model (19); ∘, ⋄, □ Theorem 3 flows 1, 2, and 3 respectively; solid lines correspond to synchronized case. Network parameters: $B = 100$ Mb, $q_{max} = 80$ packets, $\bar{T} = 20$ ms, $T_0 = 102$ ms, $T_2 = 62$ ms; approximately 0.5% bidirectional background web traffic.

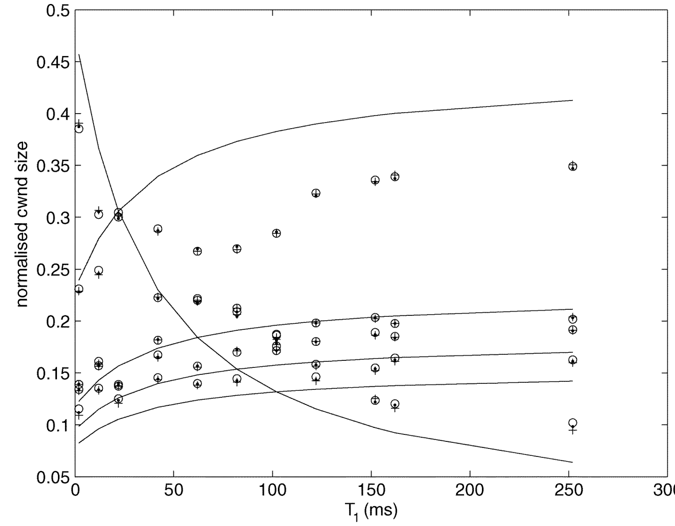

Fig. 14. Variation of mean $w_i(k)$ with propagation delay T_1 in dumbbell topology with five TCP flows. Network parameters: $B = 100$ Mb, $q_{max} = 80$ packets, $\bar{T} = 20$ ms, $T_0 = 102$ ms, $T_2 = 62$ ms; approximately 0.5% bidirectional background web traffic.

(13) gives a good approximation of these rates. However, this approximation neglects the curvature in the $cwnd$ evolution induced by time-varying round-trip time and can therefore be expected to become less accurate as the queue provisioning increases. We emphasize that the loss of predictive power is due to the validity of the approximation (13) and not the fidelity of the network model (19); a more accurate estimate of α_i would lead to better model performance. Techniques for approximating α_i when the queue is not small have already been explored in [9].

(ii) The model (19) also neglects the fact that the number of packets in flight for TCP flows is quantised: namely, restricted to integer values, owing to the packet based nature of the traffic. Hence, the accuracy of the model (19) can be expected to degrade under network conditions where the peak window size w_i of a flow is small.

B. Many Unsynchronized Flows

The foregoing results are for networks with two competing TCP sources. We note briefly that we have also validated our results against packet-level simulations for networks of up to five flows. As in the two-flow case, the simulation and analytical predictions are in close agreement; a sample of the results that we have collected is depicted in Figs. 13 and 14.

C. Limitations of Modelling Framework

The derived model (19) provides a framework for capturing the dynamics of certain types of communication networks. However, while the model encompasses features such as drop-tail queueing, flows with different round-trip times, unsynchronized loss events and the switched nature of AIMD flows, it does not capture some features of communication networks.

1) The model is only valid for a network of AIMD flows that compete for bandwidth at a single common bottleneck router with drop-tail queue, i.e., the so-called dumbbell topology.

2) The model does not include the effects of slow start and TCP timeouts (although these can easily be introduced into the model through the introduction of more matrices into the set $\mathcal{A}$).

3) In the simulation results presented, consideration is confined to situations where the queue size is small compared to the bandwidth-delay product, but this is to streamline the presentation and is not an inherent constraint of the modelling approach. There is also no assumption in the model that the queue empties following a congestion event.

4) While not intrinsic to the matrix product model itself, key assumptions for the asymptotic analysis presented are that: 1) the *pattern* of losses at each congestion event is random and independent of the congestion epoch (we do not assume that the losses seen by a flow are independent of the losses seen by other competing flows); and 2) each flow almost surely experiences a loss event provided that it sufficiently long-lived.

5) The time average results necessarily apply to long-lived flows only but our ensemble average results apply to flows of any duration.

Perhaps the most significant limitation of our model is that the probabilities of the different patterns of losses are assumed to be known beforehand (or can be measured) and are not predicted by the model. This certainly reduces the predictive power of the model. However, our objective in developing the model was not only to understand the dynamics of AIMD networks, but also to provide a basis for the design of such networks. When viewed in this context, the probabilities ρ_i play an important role in controlling the network dynamics, and as they can be controlled by the bottleneck router, are an important design parameter available to the network designer. Our model provides an analytic basis for understanding the effect of various dropping strategies of the network dynamics, and for incorporating this aspect of network dynamics into the network design procedure.

V. Related Work

An extensive literature exists relating to the modelling of TCP traffic. The well-known square-root formula of Padhye *et al.* [10] provides an approximate expression for the congestion window achieved by a TCP flow. The statistical independence assumptions in this model however neglect interactions between competing flows (e.g., the frequency of loss events is generally not independent of the values of the AIMD increase and decrease parameters of competing flows). Many of the more recent results are based on so-called fluid approaches and focus on active queueing disciplines, see, for example, [11]–[21]. Fundamental difficulties exist in applying fluid models to networks with drop-tail queues. Recently, several authors have developed new types of hybrid systems model suited to drop-tail networks, most notably Hespanha [22] and Baccelli and Hong [23]. We note that the model derived in [23] is similar to the model presented here. In particular, under mild assumptions, the sets of solutions of the model in [23] and of our model coincide, so that the results of that reference are immediately applicable to the model presented here. However, the model derived by Baccelli and Hong is both affine and the homogeneous (linear) part is characterized by general matrices (namely, not by nonnegative matrices), whereas in this paper we develop a linear, nonnegative matrix model. The properties of linearity (no affine term) and nonnegativity play a key role in the tractability of our model, both in respect of the analysis of its dynamic characteristics and of its equilibrium properties.

VI. Mathematical Derivations

Theorem 2 and Theorem 3 follow from several interesting properties of the set of matrices $\mathcal{A} = \{A_1, \dots, A_m\}$. Roughly speaking, these results may be classified as being algebraic or stochastic in nature. The purpose of this section is to elucidate these properties and to use them to prove the results given in Section III.

It was noted before that the matrices in the set $\mathcal{A}$ are not column stochastic. However, the matrices in this set are simultaneously similar to a set of column stochastic matrices under the transformation $\Gamma = \mathrm{diag}[\gamma_1, \dots, \gamma_n]$. Denote $\hat{\alpha}_j := \gamma_j \alpha_j$, $j = 1, \dots, n$. For $A \in \mathcal{A}$, determined by a choice of parameters $\beta_1(A), \dots, \beta_n(A)$ we have

$$\Gamma A \Gamma^{-1} = \begin{bmatrix} \beta_1(A) & 0 & \cdots & 0 \\ 0 & \beta_2(A) & 0 & 0 \\ \vdots & 0 & \ddots & 0 \\ 0 & 0 & \cdots & \beta_n(A) \end{bmatrix} + \frac{1}{\sum_{j=1}^n \hat{\alpha}_j} \begin{bmatrix} \hat{\alpha}_1 \\ \hat{\alpha}_2 \\ \cdots \\ \hat{\alpha}_n \end{bmatrix} [(1-\beta_1(A)), \cdots, (1-\beta_n(A))].$$

It is easy to see that the transformed matrices are column stochastic. We shall exploit this observation in the sequel as column stochastic matrices are easier to deal with than nonstochastic ones. In view of this fact we note that a Perron eigenvector of $\Gamma A_1 \Gamma^{-1}$ is given by $\bar{x}_p^T = (\hat{\alpha}_1/(1-\beta_1), \hat{\alpha}_2/(1-\beta_2), \dots, \hat{\alpha}_n/(1-\beta_n))$, and that the corresponding Perron eigenvector of A_1 is $x_p^T = (\alpha_1/(1-\beta_1), \alpha_2/(1-\beta_2), \dots, \alpha_n/(1-\beta_n))$. In the sequel we will derive results that are expressed in terms of $\bar{x}_p$. These correspond to the dynamics of the system (22) and refer to the stochastic properties of the vector $\overline{W}_\gamma(k)$. The corresponding results for the system (19) are directly deduced from these results by similarity.

A. Algebraic Properties of the Set $\mathcal{A}$

We will from now on assume without loss of generality that the matrices in the set $\mathcal{A}$ are column stochastic, which corresponds to the case $\gamma_1 = \dots = \gamma_n = 1$. Should this not be the case we can always apply the transformation Γ to obtain this property, which just amounts to a rescaling of the α_i. In the derivation of the main results of this paper we make frequent use of the fact that the matrices in the set $\mathcal{A}$, and products of matrices in this set, are nonnegative and in particular column stochastic. This observation implies the existence of an $n-1$ dimensional subspace that is invariant under $\mathcal{A}$. We will also see that the matrices in this set can be simultaneously transformed into block triangular form with an $n-1$ dimensional symmetric block. Given these observations, we will then show under mild assumptions that the distance of a matrix product of length k, constructed from matrices in $\mathcal{A}$, from the set of rank-1 matrices converges asymptotically to zero as k increases.

Lemma 1: There exists an $n-1$ dimensional subspace invariant under $\mathcal{A}$.

Proof: The row vector $v := [1, \dots, 1]$ is a left eigenvector of all of the matrices in the set $\mathcal{A}$ as they are column stochastic. This implies that the $n-1$ dimensional subspace orthogonal to v is invariant under $\mathcal{A}$ [24]. ■

Lemma 2: Consider the set of matrices $\mathcal{A}$. There exists a real nonsingular transformation T such that for all $A \in \mathcal{A}$ we have

$$T^{-1} A T = \begin{bmatrix} S & B \\ 0 & 1 \end{bmatrix} \tag{29}$$

where $S \in \mathbb{R}^{(n-1)\times(n-1)}$ is symmetric, so that in particular the eigenvalues of S are real and of absolute value ≤ 1.

Proof: We denote $\alpha = [\alpha_1 \ \dots \ \alpha_n]^T$ and $c_\alpha := (\sum_{j=1}^n \alpha_j)^{-1}$. Let $A = \Lambda + c_\alpha \alpha \beta^T \in \mathcal{A}$, where Λ is the diagonal matrix with entries equal to 1 or β_i^s and β is the corresponding vector with entries 0 or $1 - \beta_i^s$. Consider the diagonal matrix

$$D = \mathrm{diag}\left(\frac{1}{\sqrt{\alpha_1}}, \dots, \frac{1}{\sqrt{\alpha_n}}\right). \tag{30}$$

Then DAD^{-1} is a nonnegative matrix with a left eigenvector given by $z_p = vD^{-1}$, (with v defined in the proof of the previous lemma). Further, it follows that, $DAD^{-1} = \Lambda + Dc_\alpha \alpha \beta^T D^{-1}$ and by inspection $D\alpha = z_p$.

We now chose an orthogonal matrix M whose last column is $z_p/\|z_p\|$. Then e_n^T (the nth unit row vector) is a left eigenvector of $M^T DAD^{-1} M$, and furthermore

$$M^T DAD^{-1} M = M^T \Lambda M + c_\alpha M^T z_p \beta^T D^{-1} M.$$

Now as z_p is a multiple of the last column of M it follows that $M^T z_p = \|z_p\| e_n$ and hence the entries of $M^T z_p \beta^T D^{-1}$ are nonzero only in the last row. Thus, using that e_n^T is a left eigenvector we have

$$M^T DAD^{-1} M = \begin{bmatrix} S & B \\ 0 & 1 \end{bmatrix} \tag{31}$$

where $S \in \mathbb{R}^{(n-1)\times(n-1)}$ is equal to the upper left $(n-1)$-minor of $M^T \Lambda M$ and thus symmetric. The assertion follows by setting $T = D^{-1}M$. The eigenvalues of S are bounded in absolute value by 1 as the matrix A is column stochastic and thus has spectral radius equal to 1. ■

We denote the the of matrices S that appear as the upper left block in (29) by $\mathcal{S}$.

Corollary 1: Consider the system (19). Then for each $S \in \mathcal{S}$ the function $V(z((k)) = z^T(k)z(k)$ is a quadratic Lyapunov function for the dynamic system

$$\Sigma : z(k+1) = Sz(k) \tag{32}$$

i.e., for all solutions of Σ we have $V(z(k+1)) - V(z(k)) \leq 0$ for all k.

Proof: The assertion follows immediately as the matrices $\{S_1, S_2, \dots, S_m\}$ are symmetric column stochastic matrices. ■

There are some interesting consequences of Corollary 1 for products of matrices from the set $\mathcal{S}$. As these matrices are symmetric and of norm less than or equal to 1, they form what is called a *paracontracting* set of matrices. This property is defined by the requirement that

$$Sx \neq x \Leftrightarrow \|Sx\| < \|x\| \quad \forall x \in \mathbb{R}^{n-1}, S \in \mathcal{S}. \tag{33}$$

This is true for our set $\mathcal{S}$, as the matrices $S \in \mathcal{S}$ are symmetric and of spectral radius at most 1. It is known [25] that finite sets of matrices that are paracontracting have *left convergent products*, i.e., for any sequence $\{S(k)\}_{k\in\mathbb{N}}$ in $\mathcal{S}$, the following limit exists:

$$\lim_{k\to\infty} S(k)S(k-1)\dots S(0). \tag{34}$$

For related literature on paracontracting sets of matrices, refer to [25] and [26] and references therein.

In the following, we prove results on the convergence of products of the matrices in $\mathcal{A}$ to the set of column-stochastic matrices of rank 1. To this end, it will be convenient to introduce a notation that identifies each matrix $A \in \mathcal{A}$ with the sources that do not see a drop in that congestion event. Let $\mathcal{I} \subset \{1, 2, \dots, n\}$ be the index set of sources not experiencing congestion at a congestion event. (Clearly, $\mathcal{I} = \{1, 2, \dots, n\}$ can be ignored, as this means that there is no congestion.)

The matrix corresponding to an index set $\mathcal{I}$ is given by

$$A_{\mathcal{I}} = \operatorname{diag}\left(\beta_1(\mathcal{I}), \dots, \beta_n(\mathcal{I})\right) + c_\alpha \alpha \left[1-\beta_1(\mathcal{I}) \; \dots \; 1-\beta_n(\mathcal{I})\right] \tag{35}$$

where $\beta_i(\mathcal{I}) = 1$, if $i \in \mathcal{I}$ and $\beta_i(\mathcal{I}) = \beta_i^s$ otherwise and $c_\alpha := \left(\sum_{j=1}^n \alpha_j\right)^{-1}$. We now recover our set of possible matrices by

$$\mathcal{A} := \{A_{\mathcal{I}} | \mathcal{I} \subsetneq \{1, 2, \dots, n\}\} \tag{36}$$

which results in a set of $2^n - 1$ matrices, as it should. Note that all $A \in \mathcal{A}$ are column stochastic, so that they have an eigenvalue equal to 1 equal to the spectral radius.

If $\mathcal{I} \neq \emptyset$, i.e., if at least one source does not experience congestion, then the dimension of the eigenspace corresponding to 1 is equal to the number of sources not seeing the congestion event. To see this consider first the case that the first k sources $k \in \{1, \dots, n-1\}$ do not see a drop and the others do. In this case

$$A_{\{1,\dots,k\}} = \begin{bmatrix} I_{k\times k} & B \\ 0 & C \end{bmatrix} \tag{37}$$

where $B > 0$ by definition. As the matrix is column stochastic this means that all columns of C sum to a value strictly less than one, and hence $r(C) \leq \|C\|_1 < 1$ and the claim follows for $A_{\{1,\dots,k\}}$. Now an arbitrary matrix $A_{\mathcal{I}}, \mathcal{I} \neq \emptyset$ may be brought into the form (37) by permutation of the index set and we have shown the desired property.

Note also that the eigenspace of $A_{\mathcal{I}}$ associated to the eigenvalue 1 is given by

$$V_{\mathcal{I}} = \operatorname{span}\{e_i | i \in \mathcal{I}\} \tag{38}$$

where e_i denotes the ith unit vector.

Let us briefly discuss the eigenspaces of $S_{\mathcal{I}}$ corresponding to the eigenvalue 1, which we denote by $V(S_{\mathcal{I}})$. If $\mathcal{I} = \emptyset$, $\{1\}, \dots \{n\}$, then as we have seen in (38), the multiplicity of 1 as an eigenvalue of $A_{\mathcal{I}}$ is 1, so from (29) we have that $r(S_{\mathcal{I}}) < 1$. In this case we will (with slight abuse of notation) set $V(S_{\mathcal{I}}) = \{0\}$. We denote the subspace orthogonal to $[1, \dots, 1]$ by $[1, \dots, 1]^\perp$. Recall from Lemma 1, that this is an invariant subspace of $A \in \mathcal{A}$. In general, we see from (38) and Lemma 1 that if $\mathcal{I} = \{k_1, k_2, \dots, k_l\} \subsetneq \{1, \dots, n\}$ then a basis for

$$[1 \quad 1 \quad \dots \quad 1]^\perp \cap V_{\mathcal{I}} \tag{39}$$

is given, e.g., by

$$e_{k_1} - e_{k_2}, e_{k_1} - e_{k_3}, \dots, e_{k_1} - e_{k_l}.$$

Hence, the eigenspace of $V(S_{\mathcal{I}})$ is spanned by

$$M^T D^{-1}\left(e_{k_1} - e_{k_2}\right), \dots, M^T D^{-1}\left(e_{k_1} - e_{k_l}\right). \tag{40}$$

From this it follows that

$$V\left(S_{\mathcal{I}_1}\right) \cap V\left(S_{\mathcal{I}_2}\right) = V\left(S_{\mathcal{I}_1 \cap \mathcal{I}_2}\right) \tag{41}$$

justifying our abuse of notation above. In particular, $V(S_{\mathcal{I}_1}) \cap V(S_{\mathcal{I}_2}) \neq \{0\}$ if and only if $\mathcal{I}_1 \cap \mathcal{I}_2$ contains at least two elements.

Proposition 1: Let $\{S(k)\}_{k\in\mathbb{N}} \subset \mathcal{S}$ be a sequence with associated index sets $\mathcal{I}(k)$. The following statements are equivalent.

(i) For all $z_0 \in \mathbb{R}^{n-1}$, it holds that

$$\lim_{k\to\infty} S(k)S(k-1)\dots S(0)z_0 = 0.$$

(ii) For all but one $l \in \{1, \dots, n\}$, it holds that for each $k \in \mathbb{N}$, there is an $k_1 > k$ with $l \notin \mathcal{I}(k_1)$.

(iii) If $\{A_1, \dots, A_s\} \subset \mathcal{A}$ are the matrices that appear infinitely often in the sequence $A_{\mathcal{I}(k)}$, then

$$\hat{A} := \frac{1}{s}\sum_{l=1}^{s} A_l$$

is a matrix that, with the exception of at most one column, has strictly positive entries.

If in *(iii)* the kth column of $\hat{A}$ has zero entries then this column is equal to e_k.

Proof: (ii)⇔(iii): Note first that the kth column of $\hat{A}$ is not equal to e_k if and only if for one of the matrices A_l, $l = 1,\ldots,m$, the corresponding column is not a unit vector. The assumption on the matrix A_l implies that the kth source experiences a drop infinitely many times. Under assumption *(iii)* this is true for all but at most one column, which implies *(ii)*. Conversely, under the assumption *(ii)* the kth source experiences a drop infinitely many times. As there are only finitely many matrices in which the kth column is not equal to e_k, one of these appears infinitely often in the sequence of matrices and therefore in the definition of $\hat{A}$. This implies *(iii)*.

(i)⇒(ii): If *(ii)* does not hold, then (without loss of generality) there is a $s \geq 0$ such that for all $t \geq s$ we have $\{1,2\} \subset \mathcal{I}(k)$. This implies that for all $t \geq s$, the matrix $S(k)$ has the eigenspace $M^T D^{-1}(e_1 - e_2)$ as an eigenspace corresponding to the eigenvalue one. Hence, any z_0 such that $S(s)\ldots S(0)z_0$ is a multiple of $M^T D^{-1}(e_1 - e_2)$ does not satisfy *(i)*. Such a z_0 exists as all the matrices in $\mathcal{A}$ are invertible. This shows the assertion.

(ii)⇒(i): Denote $z(k) := S(k-1)\ldots S(0)z_0$. Using paracontractivity of $\mathcal{S}$ and (34) it follows that $z_\infty := \lim_{k\to\infty} z(k)$ exists. If $z_\infty = 0$ there is nothing to show. Otherwise, we claim that for some k_0 sufficiently large, it follows that $S_k z_\infty = z_\infty$ for all $k \geq k_0$. To this end, note that because $\mathcal{S}$ is finite, there exists a constant $0 < r < 1$ such that for all $S \in \mathcal{S}$ we have

$$Sz_\infty \neq z_\infty \Rightarrow \|Sz_\infty\| < r\|z_\infty\|.$$

By convergence, this implies for all $S \in \mathcal{S}$ and all k sufficiently large that

$$Sz_\infty \neq z_\infty \Rightarrow \|Sz(k)\| < \frac{1+r}{2}\|z_\infty\| < \|z_\infty\|.$$

On the other hand, the sequence $\|z(k)\|$ is decreasing, so it follows that $\|z(k)\| \geq \|z_\infty\|$ for all $k \in \mathbb{N}$. This implies that for k sufficiently large it must hold that $S(k)z_\infty = z_\infty$. This, however, means that z_∞ lies in the eigenspace $V(S(k))$ for all k large enough. From (41) it follows that at least two sources do not see a drop for all k large enough. ■

For the statement of the next result, we denote the set of column stochastic matrices of rank 1 by $\mathcal{R}$. Note that the matrices in $\mathcal{R}$ are of the form

$$\eta[1 \quad 1 \ \ldots \ 1]$$

where η is a nonnegative vector, the entries of which sum to 1. In particular, the matrices in $\mathcal{R}$ are idempotent, because $[1 \quad 1 \quad \ldots \quad 1]\eta = 1$. In the following statement, we denote the distance of a matrix B to the set $\mathcal{R}$ by

$$\operatorname{dist}(B,\mathcal{R}) := \min\left\{\|B - R\| \,|R \in \mathcal{R}\right\}.$$

Theorem 4: Let $\{A(k)\}_{k\in\mathbb{N}} \subset \mathcal{A}$ be a sequence with associated index sets $\mathcal{I}(k)$. Then each of the statements of Proposition 1 is equivalent to

$$\lim_{k\to\infty} \operatorname{dist}\left(A(k)A(k-1)\ldots A(0),\mathcal{R}\right) = 0. \qquad (42)$$

Proof: Consider Proposition 1*(ii)*. It follows from Corollary 1 that the system (22) can be transformed to an equivalent system (29). This implies that for each k the product $A(k)A(k-1)\ldots A(0)$ is similar to

$$T(k) := \begin{bmatrix} S(k)S(k-1)\ldots S(0) & * \\ 0 & 1 \end{bmatrix} \qquad (43)$$

where we do not give the expression for the entry $*$ of the matrix as it is of no relevance for our further discussion. By Proposition 1 it follows that $S(k)S(k-1)\ldots S(0) \to 0$. As the distance of $T(k)$ to a matrix of rank 1 is upper bounded by $\|S(k)S(k-1)\ldots S(0)\|$ this implies that the distance of $A(k)A(k-1)\ldots A(0)$ to the set of matrices of rank 1 converges to zero. As each of these matrices is column stochastic any limit point of the sequence $\{A(k)A(k-1)\ldots A(0)\}$ is column stochastic.

Conversely, it is clear that the (42) implies Proposition 1*(i)*. This shows the assertion. ■

The minor drawback of Proposition 1 is that no rate of convergence is supplied. Indeed, the reader may convince himself that the rate of convergence may be made arbitrarily slow by considering sequences that have repetitions of the same matrices for longer and longer intervals as $k \to \infty$. It is therefore useful to provide conditions that guarantee an exponential decay. One such condition is provided in the following proposition.

Proposition 2: For every $a \in \{n, n+1, n+2, \ldots\}$ there exists a constant $r_a < 1$ with the following property. For any sequence of index sets $\mathcal{I}(k)$ such that for all $l \in \mathbb{N}$ and all $i \in \{1,\ldots,n\}$ there is a $b \in \{la, la+1,\ldots,(l+1)a-1\}$ with $i \notin \mathcal{I}(s)$ it holds for all $k \geq k' \geq 0$ that

$$\|S(k-1)\ldots S(k')\| \leq r_k^{-2k+1} r_k^{(k-k')}$$
$$\operatorname{dist}\left(A(k-1)\ldots A(k'),\mathcal{R}\right) \leq \|T\|\,\|T^{-1}\| r_k^{-2k+1} r_k^{(k-k')}$$

with T defined by (29).

Note that any actual flow on a real network has to satisfy the assumption on the drops seen described in the previous proposition. The reason for this is that if a flow does not see a drop it will continue to increase the amount of packages sent by a constant rate. Eventually this leads to the case the the amount of packages sent exceeds the capacity of the pipe if no drops are seen. But at this point the source necessarily sees a drop. This very coarse argument shows that all realistic flows will satisfy the assumptions of the previous proposition for some k.

B. Stochastic Properties of the Set $\mathcal{A}$

We now proceed to give a number of results that relate to random products of matrices from the set $\mathcal{A}$. In this section, we assume that Assumptions 1 and 2 hold.

We first note that under our assumptions that the expectation of $\mathcal{A}$ is a positive matrix that is column stochastic with Perron eigenvector $\bar{x}_p^T = (\alpha_1/\lambda_1(1-\beta_1), \alpha_2/\lambda_2(1-\beta_2),\ldots,\alpha_n/\lambda_n(1-\beta_n))$. We then proceed to show that the expectation of

$$\Pi(k) = A(k)A(k-1)\ldots A(0)$$

is also a column stochastic matrix with the same Perron eigenvector. The second result in this section concerns the asymptotic behavior of the expectation of $\Pi(k)$. These results immediately yield Theorem 2 and Theorem 3, using the transformation Γ, if necessary.

The final results in this section revisit the convergence of $\Pi(k)$ to the set of rank-1 idempotent matrices. We show that for all $\delta > 0$ the probability of $\Pi(k)$ being at least a distance δ from the rank-1 idempotent matrices goes to zero as k becomes large.

In the following we will use the notation $A_{\mathcal{I}} = \Lambda_{\mathcal{I}} + c_\alpha \alpha \beta(\mathcal{I})^T$, where $\Lambda_{\mathcal{I}}$ denotes the diagonal matrix in (35), $c_\alpha := (\sum_{j=1}^n \alpha_j)^{-1}$ and $\beta(\mathcal{I})$ is the vector with entries $1 - \beta_i(\mathcal{I})$.

Lemma 3: Assume that $\lambda_i > 0$ for $i = 1, \ldots, n$, then the expectation

$$E(A) = \sum_{\mathcal{I}} \rho_{\mathcal{I}} A_{\mathcal{I}}$$

is positive, column stochastic, and a Perron eigenvector for it is given by

$$\bar{x}_p^T = \left(\frac{\alpha_1}{\lambda_1(1-\beta_1)}, \frac{\alpha_2}{\lambda_2(1-\beta_2)}, \ldots, \frac{\alpha_n}{\lambda_n(1-\beta_n)} \right). \quad (44)$$

Proof: By definition of the expectation and using (35) we have

$$E(A) = \sum_{\mathcal{I}} \rho_{\mathcal{I}} A_{\mathcal{I}} = \sum_{\mathcal{I}} \rho_{\mathcal{I}} \Lambda_{\mathcal{I}} + c_\alpha \alpha \sum_{\mathcal{I}} \rho_{\mathcal{I}} \beta(\mathcal{I})^T. \quad (45)$$

The ith diagonal entry of the diagonal matrix $\sum_{\mathcal{I}} \rho_{\mathcal{I}} \Lambda_{\mathcal{I}}$ is

$$\lambda_i \beta_i + (1 - \lambda_i) \quad (46)$$

and the ith entry of $\sum_{\mathcal{I}} \rho_{\mathcal{I}} \beta(\mathcal{I})$ is

$$\sum_{\mathcal{I}} \rho_{\mathcal{I}} \left(\beta(\mathcal{I})\right)_i = \lambda_i (1 - \beta_i). \quad (47)$$

Hence, the matrix $E(A)$ is of the form of A_1 defined in (21) with the same vector α and β_i replaced by $\tilde{\beta}_i := 1 - \lambda_i(1 - \beta_i) \in (0,1)$. It follows by Theorem 1 that a Perron eigenvector of $E(A)$ is given by $\bar{x}_p^T = (\alpha_1/\lambda_1(1-\beta_1), \alpha_2/\lambda_2(1-\beta_2), \ldots, \alpha_n/\lambda_n(1-\beta_n))$. ■

Lemma 4: Consider the random system (19) subject to Assumptions 1 and 2. The expectation of $\Pi(k)$ is

$$E\left(\Pi(k)\right) = E(A)^k = \left(\sum_{\mathcal{I}} \rho_{\mathcal{I}} A_{\mathcal{I}} \right)^k. \quad (48)$$

Proof: By independence, we have that the expectation of the product is the product of the expectations. This implies the equality. ■

Proof (of Theorem 2): It is sufficient to show the assertion for the case $\gamma_i = 1$, $i = 1, \ldots, n$. The assertion of (23) is shown in Lemma 4. As $E(A)$ is positive and column stochastic it follows that

$$\lim E(A)^k = \lim E\left(\Pi(k)\right) = z y^T$$

where y, z are left, respectively right, Perron eigenvectors of $E(A)$. As $E(A)$ is column stochastic we may normalize $y = y_P^T = [1 \ \ldots \ 1]$. Finally, the assertion concerning x_P follows from Lemma 3. ■

Proposition 3: Consider the random system (19) subject to Assumptions 1 and 2. Then, with probability one,

$$\lim_{k \to \infty} \operatorname{dist}\left(A(k)A(k-1)\ldots A(0), \mathcal{R}\right) = 0.$$

Proof: Under the assumptions that the λ_j are positive and the independence assumptions, with probability one each source will see infinitely many drops. Now the result follows from Theorem 4. ■

C. Proof of Theorem 3

We now proceed to present an outline of the proof of main result of this paper, Theorem 3. In [27], it is shown that the result can be derived from general results on Markov e-chains. The technical preparations that this line of argumentation requires, however, are beyond the scope of this paper. In [8], we give a proof that relies on fairly elementary arguments in order to keep the main ideas accessible.

Outline of Proof: We are interested in the asymptotic behavior of the average window variable $\overline{W}(k)$:

$$\overline{W}(k) = \frac{1}{k} \sum_{i=0}^{k-1} W1) = \frac{1}{k} \left(\Pi(k-1) + \ldots + \Pi(0)\right) W(0)$$

as k tends to infinity. Our proof consists of the following main steps.

Step 1) For a fixed k_0 we partition each sufficiently long product $\Pi(k) = \Phi(k)\Psi(k)$, where $\Phi(k)$ is the leading product of length k_0. We know that as $k_0 \to \infty$, the leading product approaches almost surely the set of rank one matrices, which implies that $\Phi(k)\Psi(k) \approx \Phi(k)$ as all matrices involved are column stochastic.

Step 2) We thus may approximate $\overline{W}(k)$ as

$$\begin{aligned} \overline{W}(k) &= \frac{1}{k+1} \left(\Pi(k) + \ldots + \Pi(0)\right) W(0) \\ &= \frac{1}{k+1} \left(R(k) + \Delta(k) + \ldots + R(l) + \Delta(l) \right. \\ &\qquad \left. + \Pi(l-1) + \ldots + \Pi(0)\right) W(0) \quad (49) \\ &\approx \frac{1}{k+1} \left(R(k) + \ldots + R(l) + \Pi(l-1) \right. \\ &\qquad \left. + \ldots + \Pi(0)\right) W(0) \quad (50) \end{aligned}$$

where $R(k), \ldots, R(l)$ are column stochastic rank-1 matrices and $\Delta(k), \ldots, \Delta(l)$ are error terms that approach 0 as $k_0 \to \infty$.

Step 3) Using the law of large numbers, it is then seen that $1/k(R(k) + \ldots + R(l))$ can be approximated as $(\sum_{i=1}^m \rho_i A_i)^l$.

Step 4) And it follows that

$$\lim_{k \to \infty} \overline{W}(k) = x_p \bar{y}_p^T W(0)$$

where $\bar{y}_p^T = (1, \ldots, 1)$.

VII. Conclusion

In this paper, we have presented and validated using packet level simulations, a random matrix model that describes the dynamic behavior of a network of n AIMD flows that compete for shared bandwidth via a bottleneck router employing drop-tail queueing. We have used this model to relate several important network properties to properties of sets of nonnegative matrices that arise in the study of such networks. We have also derived under simplifying assumptions a number of analytic results that characterize the asymptotic time-average and ensemble-average throughput of such networks.

References

[1] R. N. Shorten, D. J. Leith, J. Foy, and R. Kilduff, "Analysis and design of AIMD congestion control algorithms in communication networks," *Automatica*, vol. 41, pp. 725–730, 2005.

[2] A. Berman and R. Plemmons*, Nonnegative Matrices in the Mathematical Sciences*. Philadelphia, PA: SIAM, 1979.

[3] R. Shorten, D. Leith, J. Foy, and R. Kilduff, "Analysis and design of synchronized communication networks," in *Proc. 12th Yale Workshop on Adaptive and Learning Systems*, 2003.

[4] A. Berman, R. Shorten, and D. Leith, "Positive matrices associated with synchronized communication networks," *Linear Algebra Appl.*, vol. 393, pp. 47–54, 2004.

[5] L. Xu, K. Harfoush, and I. Rhee, "Binary increase congestion control (BIC) for fast long-distance networks," in *Proc. IEEE INFOCOM*, Hong Kong, 2004, pp. 2514–2524.

[6] S. Floyd and V. Jacobson, "Traffic phase effects in packet-switched gateways," *J. Internetworking: Practice and Experience* vol. 3, no. 3, pp. 115–156, Sep. 1992 [Online]. Available: citeseer.ist.psu.edu/floyd92traffic.html

[7] W. Willinger, M. S. Taqqu, R. Sherman, and D. V. Wilson, "Self-similarity through high-variability: statistical analysis of Ethernet LAN traffic at the source level," *IEEE/ACM Trans. Netw.*, vol. 5, no. 1, pp. 71–86, Feb. 1997.

[8] R. Shorten, F. Wirth, and D. Leith, Positive Matrices and the Internet Hamilton Institute Tech. Rep. [Online]. Available: http://www.hamilton.ie

[9] D. Leith and R. Shorten, "Modelling TCP throughput and fairness," Proc. of Networking 2004 Lecture Notes in Computer Science, vol. 3042. Berlin, Germany, Springer-Verlag, 2004, pp. 938–948.

[10] J. Padhye, V. Firoiu, D. F. Towsley, and J. F. Kurose, "Modeling TCP Reno performance: a simple model and its empirical validation," *IEEE/ACM Trans. Netw.*, vol. 8, no. 2, pp. 133–145, Apr. 2000.

[11] F. P. Kelly, "Mathematical modelling of the internet," in *Proc. ICIAM'99, 4th Int. Congr. Industrial Applied Mathematics*, Edinburgh, 2000, pp. 105–116.

[12] S. Low, F. Paganini, and J. Doyle, "Internet congestion control," *IEEE Contr. Syst. Mag.*, vol. 32, no. 1, pp. 28–43, 2002.

[13] S. Mascolo, "Congestion control in high speed communication networks using the Smith principle," *Automatica*, vol. 35, pp. 1921–1935, 1999.

[14] S. S. Kunniyur and R. Srikant, "Stable, scalable, fair congestion control and AQM schemes that achieve high utilisation in the internet," *IEEE Trans. Automat. Contr.*, vol. 48, no. 11, pp. 2024–2029, 2003.

[15] L. Massoulie, "Stability of distributed congestion control with heterogeneous feedback delays," *IEEE Trans. Automat. Contr.*, vol. 47, no. 6, pp. 895–902, 2002.

[16] C. V. Hollot, V. Misra, D. Towsley, and W. Gong, "Analysis and design of controllers for AQM routers supporting TCP flows," *IEEE Trans. Automat. Contr.*, vol. 47, no. 6, pp. 945–959, 2002.

[17] G. Vinnicombe, "On the stability of networks operating TCP-like congestion control," in *Proc. 15th IFAC World Congr. Automatic Control*, Barcelona, Spain, Jul. 2002.

[18] R. Johari and D. Tan, "End-to end congestion control for the internet: delays and stability," *IEEE/ACM Trans. Netw.*, vol. 9, no. 6, pp. 818–832, Dec. 2001.

[19] Y. Chait, C. V. Hollot, V. Misra, H. Han, and Y. Halevi, "Dynamic analysis of congested TCP networks," in *Proc. American Control Conf.*, San Diego, CA, Jun. 1999, pp. 2430–2435.

[20] C. V. Hollot and Y. Chait, "Nonlinear stability analysis of a class of TCP/AQM networks," in *Proc. 40th IEEE Conf. Decision and Control*, Orlando, FL, Dec. 2001, pp. 2309–2314.

[21] C. V. Hollot, V. Misra, D. Towsley, and W. Gong, "A control theoretic analysis of RED," in *Proc. IEEE INFOCOM*, Anchorage, AK, Apr. 2001, pp. 1510–1519.

[22] J. Hespanha, "Stochastic hybrid systems: application to communication networks," in *Proc. Hybrid Systems, Computation and Control (HSCC2004)*, Mar. 2004, pp. 387–401.

[23] F. Baccelli and D. Hong, "AIMD, fairness and fractal scaling of TCP traffic," in *Proc. IEEE INFOCOM*, New York, Jun. 2002, pp. 229–238.

[24] R. A. Horn and C. R. Johnson*, Matrix Analysis*. Cambridge, U.K.: Cambridge Univ. Press, 1985.

[25] A. Vladimirov, L. Elsner, and W.-J. Beyn, "Stability and paracontractivity of discrete linear inclusions," *Linear Algebra Appl.*, vol. 312, no. 1–3, pp. 125–134, 2000.

[26] W.-J. Beyn and L. Elsner, "Infinite products and paracontracting matrices," *Electron. J. Linear Algebra*, vol. 2, pp. 1–8, 1997.

[27] F. Wirth, R. Stanojević, R. Shorten, and D. Leith, "Stochastic equilibria of AIMD communication networks," *SIAM J. Matrix Analysis Appl.*, 2006.

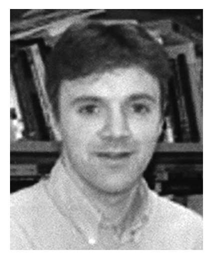

Robert Shorten graduated from the University College Dublin (UCD), Ireland, in 1990 with a First Class Honours B.E. degree in electronic engineering.

From 1993 to 1996, he was the holder of a Marie Curie Fellowship to conduct research at the Daimler-Benz Research Institute for Information Technology in Berlin. In 1997, he was awarded a European Presidential fellowship to return to Ireland. He is a co-founder and a Senior Researcher of the Hamilton Institute, NUI Maynooth, Ireland, and is an Editor of the *IEE Proceedings on Control Theory*. His research interests include stability theory, linear algebra, and network congestion control.

Fabian Wirth received the Diploma, Dr. rer. nat. and venia legendi in mathematics from the University of Bremen, Germany, where he was with the Institute for Dynamical Systems and the Centre for Technomathematics.

He is a Senior Researcher at the Hamilton Institute, NUI Maynooth, Ireland, where he works on the dynamics of communication networks. His interests include stability theory of dynamical systems and robust stability.

Douglas Leith graduated from the University of Glasgow, Scotland, U.K., in 1986 with a first class B.Sc. (Eng.) degree in electronics and electrical engineering and computer science and received the Ph.D. degree, also from the University of Glasgow, in 1989.

Following the award of a Royal Society personal research fellowship to study nonlinear control, in 2001 Prof. Leith joined the National University of Ireland, Maynooth, as Director of the Hamilton Institute. His current research interests include internet congestion control and dynamics, resource allocation in wireless networks, and nonlinear time series analysis.